arXiv:1703.02660v1 [cs.LG] 8 Mar 2017

Towards Generalization and Simplicity in Continuous Control

Aravind Rajeswaran* **Kendall Lowrey*** **Emanuel Todorov** **Sham Kakade**

Abstract

This work shows that policies with simple linear and RBF parameterizations can be trained to solve a variety of continuous control tasks, including the OpenAI gym benchmarks. The performance of these trained policies are competitive with state of the art results, obtained with more elaborate parameterizations such as fully connected neural networks. Furthermore, existing training and testing scenarios are shown to be very limited and prone to over-fitting, thus giving rise to only trajectory centric policies. Training with a diverse initial state distribution is shown to produce more global policies with better generalization. This allows for interactive control scenarios where the system recovers from large on-line perturbations; as shown in the video. [1]

1. Introduction

Deep learning has recently achieved impressive results on a number of hard problems, including sequential decision making in game domains (Mnih, 2015; Silver, 2016). This has motivated efforts to extend and adapt such techniques to the control of physical systems. Neural network controllers trained with Reinforcement Learning have recently produced rich motor behaviors on a number of simulated control tasks (Lillicrap et al., 2015; Schulman et al., 2016). The complexity of the systems being controlled is not yet at the level of what can be achieved through trajectory optimization in simulation (Tassa et al., 2012; Mordatch et al., 2012; Al Borno et al., 2013) or with hand-crafted controllers on physical robots (e.g. Boston Dynamics). However, RL approaches are exciting since in principle they are generic, model-free, and highly automatic.

Brockman et al. (2016) released a suite of physics models and associated control tasks implemented in the MuJoCo simulator (Todorov et al., 2012). This has provided common benchmarks which have accelerated research and enabled objective comparisons. Duan et al. (2016) surveyed various algorithms and studied their performance on these benchmark tasks. The best performing algorithms seem to draw upon the benefits of the Natural Gradient (Kakade, 2001) and expressive power of neural networks.

While the training algorithms and resulting performance have been well studied, a clear understanding of design choices like policy architecture is lacking. Recent research effort in other areas of AI such as computer vision have been devoted to finding tailored architectures for the respective problems (He et al., 2016). However, for control tasks, recent works have focused almost exclusively on small fully connected neural networks and occasionally recurrent networks (see Duan et al. (2016) for a survey). As a first step towards identifying tailored architectures for control, this work considers simple linear and RBF policies.

This work empirically explores the issue of architecture choice for benchmark control problems, and finds that the research area is much more open than the recent literature might suggest. In particular, the growing sentiment in the community that (i) significant improvements on common benchmark control tasks are due to the use of neural network policies; and (ii) the existing benchmarks provide a sufficient basis for future research, appears to be inaccurate. This work shows that simpler architectures – linear and RBF policies – are competitive with state of the art results on continuous control tasks. The work also shows that existing task designs are insufficient to test and understand generalization capabilities of trained policies. For example, always initializing a walking agent upright does not enable the agent to learn how to get up if it falls down, since it overfits to specific regions of the state space. Combined, these observations suggest that trajectory centric policies are sufficient to succeed as per the gym-v1 task definitions, and that simple representations have adequate capacity for such policies. In order to promote generalization and training of global policies for the walking task, a more diverse initial state distribution can be used. In this way the same policy can learn to get up from the ground, locomote, as well as recover from a wide range of perturbations. The attached video illustrates simulations where the agent recovers from on-line perturbations. Interestingly, the RBF policy is able to learn a global policy to solve this modi-

* Equal contributions. All authors are with the University of Washington, Department of Computer Science & Engineering.
This work was supported by the US National Science Foundation.
Correspond to: <aravraj> at <cs.washington.edu>

[1] Project site: https://sites.google.com/view/simple-pol

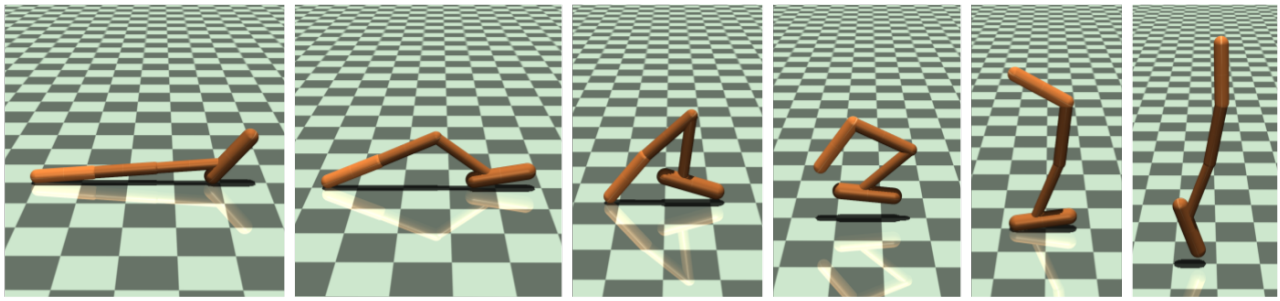

Figure 1. Hopper completes a get-up sequence before moving to its normal forward walking behavior. The getup sequence is learned along side the forward hopping through our new tasks formulations.

fied task as well. This suggests that expressive power of policies may not be the bottle neck for most of the gym-v1 continuous control tasks.

This work does not attempt to provide a definitive answer in terms of the ideal architecture choice for control. Rather, the results herein suggest that the current set of benchmark tasks are insufficient to provide insights to this question. In addition, this work suggests that a rethinking about the problem setting is required to understand robustness and generalization capabilities, as well as potential application in robotics. As a first step towards this goal, this work considers interactive variants of the tasks in the gym suite. Here, at test time, a user is allowed to interactively apply perturbations such as pushing or rotating the agent. The agent should be able to recover from these perturbations and succeed in its objective. We emphasize that these interactive tasks are not an ultimate test for generalization capability, but rather a step towards this goal.

Ideally, a controlled scientific study would seek to isolate the challenges related to architecture, task design, and training method for separate study. In practice, this is not feasible as the results are coupled with the training methods (and there is little understanding of this interplay). For simplicity, this work utilizes a straightforward natural policy gradient method with normalized stepsize, which is closely related to the TRPO method (Schulman et al., 2015).

2. Problem Formulation

We consider Markov Decision Processes (MDPs) in the average reward setting, which is defined using the tuple: $\mathcal{M} = \{\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{T}, \rho_0\}$. $\mathcal{S} \subseteq \mathbb{R}^n$, $\mathcal{A} \subseteq \mathbb{R}^m$, and $\mathcal{R} : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ are (continuous) set of states, set of actions, and reward function respectively, and have the usual meaning. $\mathcal{T} : \mathcal{S} \times \mathcal{A} \to \mathcal{S}$ is the stochastic transition function and ρ_0 is the probability distribution over initial states. We wish to solve for a stochastic policy of the form $\pi : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$, which optimizes the objective function:

$$\eta(\pi) = \lim_{T \to \infty} \frac{1}{T} \mathbb{E}_{\pi, \mathcal{M}} \left[\sum_{t=1}^{T} r_t \right]. \tag{1}$$

Since we use simulations with finite length rollouts to estimate the objective and gradient, we approximate $\eta(\pi)$ using a finite T. In this finite horizon rollout setting, we define the value, Q, and advantage functions as follows:

$$V^\pi(s,t) = \mathbb{E}_{\pi, \mathcal{M}} \left[\sum_{t'=t}^{T} r_{t'} \right]$$

$$Q^\pi(s,a,t) = \mathbb{E}_{\mathcal{M}} \Big[\mathcal{R}(s,a) \Big] + \mathbb{E}_{s' \sim \mathcal{T}(s,a)} \Big[V^\pi(s', t+1) \Big]$$

$$A^\pi(s,a,t) = Q^\pi(s,a,t) - V^\pi(s,t)$$

Note that even though the value functions are time-varying, we still optimize for a stationary policy. We consider parametrized policies π_θ, and hence wish to optimize for the parameters (θ). Thus, we overload notation and use $\eta(\pi)$ and $\eta(\theta)$ interchangeably.

3. Algorithm

Using the likelihood ratio approach and Markov property of the problem, the policy gradient is derived to be:

$$\nabla_\theta \hat{\eta}(\theta) = g = \frac{1}{T} \sum_{t=0}^{T} \nabla_\theta \log \pi_\theta(a_t | s_t) \hat{A}^\pi(s_t, a_t, t) \tag{2}$$

Gradient ascent using this "vanilla" gradient is sub-optimal since it is not the steepest ascent direction in the metric of the parameter space (Amari, 1998; Kakade, 2001). The steepest ascent direction is obtained by solving the following local optimization problem around iterate θ_k:

$$\begin{aligned} \underset{\theta}{\text{maximize}} \quad & g^T(\theta - \theta_k) \\ \text{subject to} \quad & (\theta - \theta_k)^T F_{\theta_k} (\theta - \theta_k) \leq \delta, \end{aligned} \tag{3}$$

where F_{θ_k} is the Fisher Information Metric at the current iterate θ_k. We estimate this metric as

$$\hat{F}_{\theta_k} = \frac{1}{T}\sum_{t=0}^{T} \nabla_\theta \log \pi_\theta(a_t|s_t) \nabla_\theta \log \pi_\theta(a_t|s_t)^T. \quad (4)$$

This yields the steepest ascent direction to be $\hat{F}_{\theta_k}^{-1} g$ and corresponding update rule: $\theta_{k+1} = \theta_k + \alpha \hat{F}_{\theta_k}^{-1} g$. Here α is the step-size or learning rate parameter. Empirically, we observed that choosing a fixed value for α or an appropriate schedule is difficult. Thus, we use the normalized gradient ascent procedure, where the normalization is under the Fisher metric. This procedure can be viewed as picking a normalized step size δ as opposed to α, and solving the optimization problem in (3). This results in the following update rule:

$$\theta_{k+1} = \theta_k + \sqrt{\frac{\delta}{g^T \hat{F}_{\theta_k}^{-1} g}}\, \hat{F}_{\theta_k}^{-1} g. \quad (5)$$

A dimensional analysis of these quantities reveal that α has the unit of return^{-1} whereas δ is dimensionless. Though units of α are consistent with a general optimization setting where step-size has units of objective^{-1}, in these problems, picking a good α that is consistent with the scales of the reward was difficult. On the other hand, a constant normalized step size was numerically more stable and easier to tune: for all the results reported in this paper, the same $\delta = 0.05$ was used. Henceforth, we use the notation $\alpha_n = \sqrt{\delta / g^T \hat{F}_{\theta_k}^{-1} g}$. When more than one trajectory rollout is used per update, the above estimators can be used with an additional averaging over the trajectories.

For estimating the advantage function, we use the GAE procedure (Schulman et al., 2016). This requires learning a function that approximates V_k^π, which is used to compute A_k^π along trajectories for the update in (5). GAE helps with variance reduction at the cost of introducing bias, and requires tuning hyperparameters like a discount factor and an exponential averaging term. Good heuristics for these parameters have been suggested in prior work. The same batch of trajectories cannot be used for both fitting the value function baseline, as well as estimating g using (2), since it will lead to overfitting and a biased estimate. Thus, we use the trajectories from iteration $k-1$ to fit the value function, essentially approximating V_{k-1}^π, and use trajectories from iteration k for computing A_k^π and g. Similar procedures have been adopted in prior work (Duan et al., 2016).

The optimization procedure we use is closely related to TRPO, but with a few differences. In TRPO, the Fisher metric is estimated as $\hat{F}_{\theta_k}^{\text{TRPO}} = \nabla_\theta^2 \bar{KL}\big(\pi(\cdot;\theta_k)||\pi(\cdot;\theta)\big)$, which incurs finite sample errors. TRPO also avoids computing $\hat{F}_{\theta_k}^{-1}$ fully by using a conjugate gradient procedure and Hessian-vector products. We also employ a similar procedure for computing the Natural gradient. First, we recognize that $\hat{F}_{\theta_k} = \frac{1}{N}\sum_{i=1}^{N} \nabla_\theta L_i \nabla_\theta L_i^T$, where $\nabla_\theta L_i$ is short hand for $\nabla_\theta \log \pi_\theta(a_i|s_i)$ for a sample (s_i, a_i) collected in iteration k with policy $\pi(\cdot, \theta_k)$. With this, efficient matrix-vector products (i.e. $\hat{F}_{\theta_k} v$) can be computed as $\frac{1}{N}\sum_{i=1}^{N} \nabla_\theta L_i \Big(\nabla_\theta L_i^T v\Big)$ for a given vector v. This avoids explicit computation of $\hat{F}_{\theta_k}$ or $\hat{F}_{\theta_k}^{-1}$, which could be costly when optimizing many parameters. Due to the use of Conjugate Gradient for computing $\hat{F}_{\theta_k}^{-1} g$, this procedure is also sometimes called Truncated Natural Gradient. In addition, TRPO constrains the KL divergence between successive policy iterates. The constraint in (3) can be interpreted as an approximation to the KL-divergence. In practice, we found that the KL constraint is violated only a very small fraction of times during the training run, suggesting that with a good approximation of the Fisher metric, a hard bound on the KL divergence may not be required.

Algorithm 1 Policy Search with Natural Gradient

1: Initialize policy parameters to θ_0
2: **for** $k = 1$ **to** K **do**
3: Collect trajectories $\{\tau^{(1)}, \ldots \tau^{(N)}\}$ by rolling out the stochastic policy $\pi(\cdot; \theta_k)$.
4: Compute $\nabla_\theta \log \pi(a_t|s_t;\theta_k)$ for each (s, a) pair along trajectories sampled in iteration k.
5: Compute advantages A_k^π based on trajectories in iteration k and approximate value function V_{k-1}^π.
6: Compute policy gradient according to (2).
7: Compute the Fisher matrix (4) and perform gradient ascent (5).
8: Update parameters of value function in order to approximate $V_k^\pi(s_t^{(n)}) \approx R(s_t^{(n)})$, where $R(s_t^{(n)})$ is the empirical return computed as $R(s_t^{(n)}) = \sum_{t'=t}^{T} \gamma^{(t'-t)} r_t^{(n)}$. Here n indexes over the trajectories.
9: **end for**

3.1. Policy Architecture

We first consider a linear policy that directly maps from the observations to the motor torques. We use the same observations as specified in the Gym-v1 tasks which include joint positions, joint velocities, and for some tasks, information related to contacts. Thus, the policy mapping is $a_t \sim \mathcal{N}(W s_t + b, \sigma)$, and the goal is to learn W, b, and σ. For most of these tasks, the observations correspond to the state of the problem (in relative coordinates). Thus, we use the term states and observations interchangeably. In general, the policy is defined with observations as the input, and hence is trying to solve a POMDP.

Secondly, we consider a more expressive policy based on Fourier features of the observations. Since these fea-

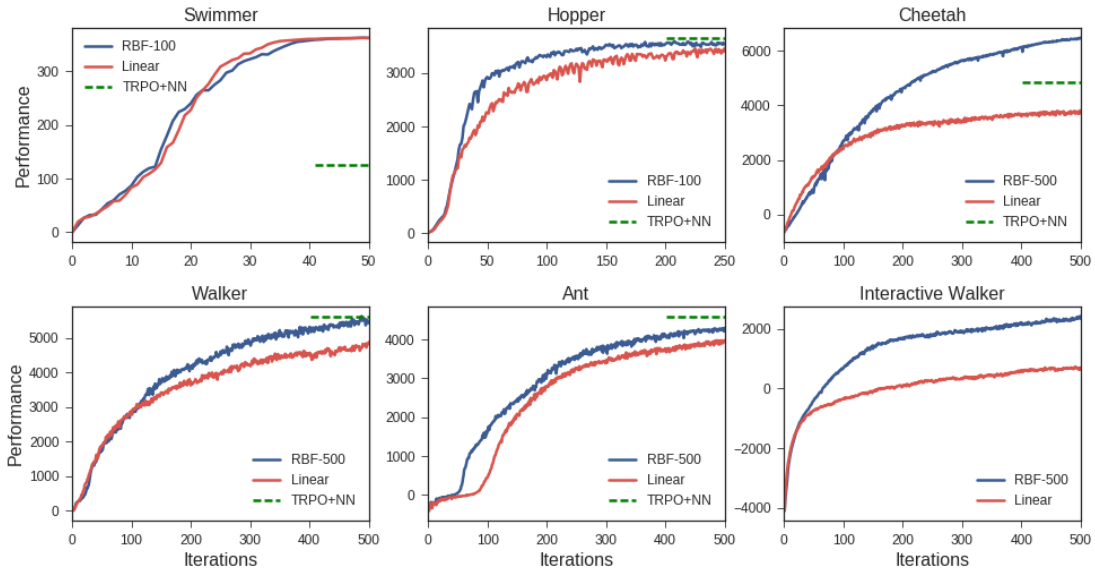

Figure 2. Learning curves for the Linear and RBF policy architectures. The green line corresponding to the reward achieved by neural network policies on the OpenAI Gym website, as of 02/24/2017 (trained with TRPO). It is observed that for all the tasks, linear and RBF parameterizations are competitive with state of the art results. The learning curves depicted are for the stochastic policies, where the actions are sampled as $a_t \sim \pi_\theta(s_t)$. The learning curves have been averaged across three runs with different random seeds.

tures approximate the RKHS features under an RBF Kernel (Rahimi & Recht, 2007), we call this policy parametrization the RBF policy. These features are constructed as:

$$y_t^{(i)} = \sin\left(\frac{\sum_j P_{ij} s_t^{(j)}}{\nu} + \phi^{(i)}\right), \tag{6}$$

where each element P_{ij} is drawn from $\mathcal{N}(0,1)$, ν is a bandwidth parameter chosen approximately as the average pairwise distances between different observation vectors, and ϕ is a random phase shift drawn from $U[-\pi,\pi)$. Thus the policy is $a_t \sim \mathcal{N}(Wy_t+b,\sigma)$, where W, b, and σ are trainable parameters. This architecture can also be interpreted as a two layer neural network: the bottom layer with fixed random weights and sinusoidal activation and the top with learned weights and linear activation.

4. Experiments and Results

4.1. Gym-v1 tasks

We test the algorithm from Section 3 on the OpenAI gym-v1 benchmarks simulated in MuJoCo (Todorov et al., 2012). The set of tasks include the swimmer, hopper, half-cheetah, walker, and ant locomotion tasks. The experimental goal is to understand the performance level achievable with simple linear policies, and whether this performance can be further improved with the RBF policy.

Table 1. Final performances of the policies

Task	Linear		RBF		NN
	stoc	mean	stoc	mean	TRPO
Swimmer	362	**366**	361	365	131
Hopper	3466	3651	3590	**3810**	3668
Cheetah	3810	4149	6477	**6620**	4800
Walker	4881	5234	5631	**5867**	5594
Ant	3980	4607	4297	4816	**5007**

Table 2. Number of episodes required to achieve threshold

Task	Th.	Linear	RBF	TRPO+NN
Swimmer	325	**1450**	1550	N-A
Hopper	3120	13920	**8640**	10000
Cheetah	3430	11250	6000	**4250**
Walker	4390	36840	25680	**14250**
Ant	3580	39240	**30000**	73500

Figure 2 presents the learning curves for the Gym-v1 tasks along with the performance levels reported in prior work using TRPO and neural network policies. Table 1 also summarizes the final scores, where "stoc" refers to the stochastic policy with actions sampled as $a_t \sim \pi_\theta(s_t)$, while "mean" refers to use of mean of the policy, with actions computed as $a_t = \mathbb{E}[\pi_\theta(s_t)]$. We see that the linear policy is competitive on most tasks, while the RBF policy can outperform previous results on four of the five considered tasks. Though we were able to train neural network policies that match the results reported in literature, we have used results reported in prior work for an objective comparison. Visualizations of the trained linear and RBF policies are presented in the supplementary video. Given the simplicity of these policies, it is surprising that they can produce such elaborate behaviors.

Table 2 presents the number of samples needed for the policy performance to reach a threshold value for reward. The threshold value is computed as 90% of the final score achieved by the stochastic linear policy. We visually verified that policies with these scores are proficient at the task, and hence the chosen values correspond to meaningful performance thresholds. We see that linear and RBF policies are able to learn faster on three of the five tasks.

This work primarily considers comparison with neural network policies trained with TRPO, since it has demonstrated impressive results and is closest to the algorithm used in this work. We defer comparisons to policies trained with value function based methods or actor critic methods, such as DDPG (Lillicrap et al., 2015) for future work.

An interesting observation relates to the number of parameters and training time required for the different policy architectures. As an example, for the hopping task, the linear policy requires a total of 39 parameters to be optimized, in contrast to prior work which optimizes for 100x as many. In wall-clock times, it was possible to train policies (both linear and RBF) for all the tasks in under an hour, with the simpler tasks like swimmer and hopper requiring under 10 minutes. In contrast, more elaborate parameterizations require multiple hours of training.

The Gym-v1 tasks were designed in such a way that, in order to succeed, the policy needs to output correct actions only for a small subset of meaningful states. Thus, the representational capacity needed is minimal, which enables the linear policy to be competitive. In order to encourage learning of more global policies and to push the representational limits, we test the performance of these policy architectures on more challenging variants of these tasks in the next section. Finding global policies on these tasks with on-line model based optimization has been a challenging benchmark in literature due to under-actuation, high dimensionality, and presence of contacts (Erez et al., 2011).

Table 3. Modified Tasks

Task	Model Change	Initialization
Swimmer	-	Starts in any Z-orientation
Hopper	Gearing	Can start prone
Cheetah	Mass	-
Walker	Gearing, Mass	Can start prone
Ant	Gearing, Mass	Starts in any Z-orientation

4.2. Modified Tasks

To push the representational capabilities of the policy architectures, we modified the Gym-v1 tasks as follows:

1. Removing termination conditions to allow agent access to potentially all parts of the state space.
2. Wider initial state distribution to force generalization.
3. Reward shaping and more realistic choice of physical properties (e.g. mass, friction) to achieve more natural looking behaviors such as walking gaits.

Combined, these modifications require that the learned policies not only make progress towards maximizing the reward, but also recover from adverse conditions and resist perturbations. An example of this is illustrated in Figure 1, where the hopper model demonstrates a get-up sequence to continue forward movement.

A short summary and high-level objective of the modified tasks are presented below and in Table 3:

1. Swimmer and Ant: heading angle initialized randomly in $[-\pi/2, \pi/2]$ relative to the intended travel direction.
2. Hopper and Walker: In the Gym-v1 tasks, the agent is always initialized upright and episode is terminated when the center of mass falls below a height threshold. We consider two variants: (a) no termination, but the agent is always initialized upright; (b) no termination, and the agent is initialized from a wider distribution, such as laying on the ground for half the episodes.

Termination conditions help with reward shaping by penalizing sequences of actions that bring the agent to undesirable parts of the state space. From an MDP perspective, termination corresponds to transitioning to a pseudo absorbing state when the system moves outside the allowable state space. Removing termination allows the agent to both enter, and hopefully, leave the undesirable states, the goal being for additional robustness to unexpected consequences. To further encourage this goal, we induce exploration of unfavorable states by initializing our agents in these undesirable configurations. In order to achieve a high score, the agent must be able to learn and represent all the necessary skills, such as turning in case of swimmer, and getting up from the ground for hopper and walker.

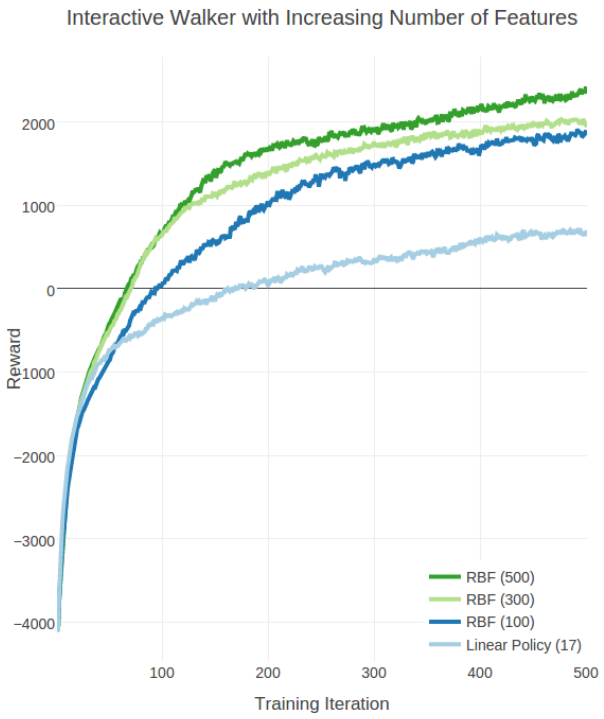

Figure 3. We train a policy on our interactive 2-D Walker task with varying numbers of Fourier features. A linear policy is included as a baseline. The linear policy is able to make reasonable forward progress, but has trouble getting up from the prone position. Larger policies are able to learn behaviors that are able to generate greater rewards, most likely by recovering from prone positions more quickly. Shown above is the learning curve corresponding to stochastic policy averaged over five runs with different random seeds.

In addition, termination conditions also interfere with the optimization procedure, especially if it is possible to enter the absorbing terminal state from the initial state through some actions. In such situations, two policies that produce similar actions (say, have small KL-divergence), can have very different evaluations. The noise in the policy could be driving the system away from these absorbing states, thereby adversely affecting the smoothness of the optimization landscape. This problem is further exacerbated when we consider the mean of the stochastic policy, which does not have the necessary noise to escape the absorbing states.

We first study the representational capacity of policies on the modified walker task in Figure 3. Increasing the Fourier features allows for more expressive policies and consequently allow for achieving a higher score. The policy with 500 RBF features out performs policies with only 300 and 100 features. The linear policy also makes forward progress but is unable to learn as efficient a walking gait.

We also test the robustness of our policies in these new environments by perturbing the system with an external force. This external force represents an unforeseen change which the agent has to resist or overcome. In these tasks, perturbations are not applied to the system during the training phase. Thus, the ability to generalize and resist perturbations come entirely out of the states visited by the agent during training. Figure 4 shows that a more expressive policy is able to better perform in this setting, and also that diverse initializations are important to obtain the best results. For example, in the situation where the simulated robot is always initialized upright (but without any termination conditions), the agent learns to recover from perturbations that make it fall forward, but it cannot recover if it falls over backwards–the agent is more likely to fall forward during training when learning to propel itself forward. Similarly, for the swimmer task, the policy is unable to generalize to various heading angles. This indicates that careful design of initial state distributions are crucial for generalization, and to enable the agent to learn a wide range of skills. Without diverse initializations and if termination conditions are used, the learned policy is effectively only trajectory centric, albeit with potentially larger coverage that those obtained using trajectory optimization methods like iLQG (Tassa et al., 2012).

5. Discussion and Future Work

5.1. Sample Complexity and Optimization

Even though linear policies were able to produce surprisingly good results in Gym-v1 tasks, and perform well even in the modified tasks, they do not appear to be favorable in terms of sample complexity when optimized with policy gradient methods. As can be seen from the learning curves in Figure 2 as well as Figure 3, the linear policy is able to learn faster initially, but this benefit quickly diminishes after a few dozen iterations. Thus, over-parametrized representations appear to be beneficial from an optimization perspective, but the ability of these over-parametrized representations (neural networks, RBFs etc.) to generalize is yet to be understood completely, and is beyond the scope of this work. Understanding these issues likely requires a rethinking of the problem settings and more appropriate benchmark tasks. In addition, it would be interesting to see if simple policies can be optimized more efficiently using alternate direct policy search methods such as CMA-ES, and value function based methods.

Based on the observation that linear and RBF policies work remarkably well on the Gym-v1 tasks, we chose these architectures for the interactive tasks we developed. A careful study of the performance with different architectures including fully connected and recurrent networks remains to be done on these tasks, and would make for interesting future work.

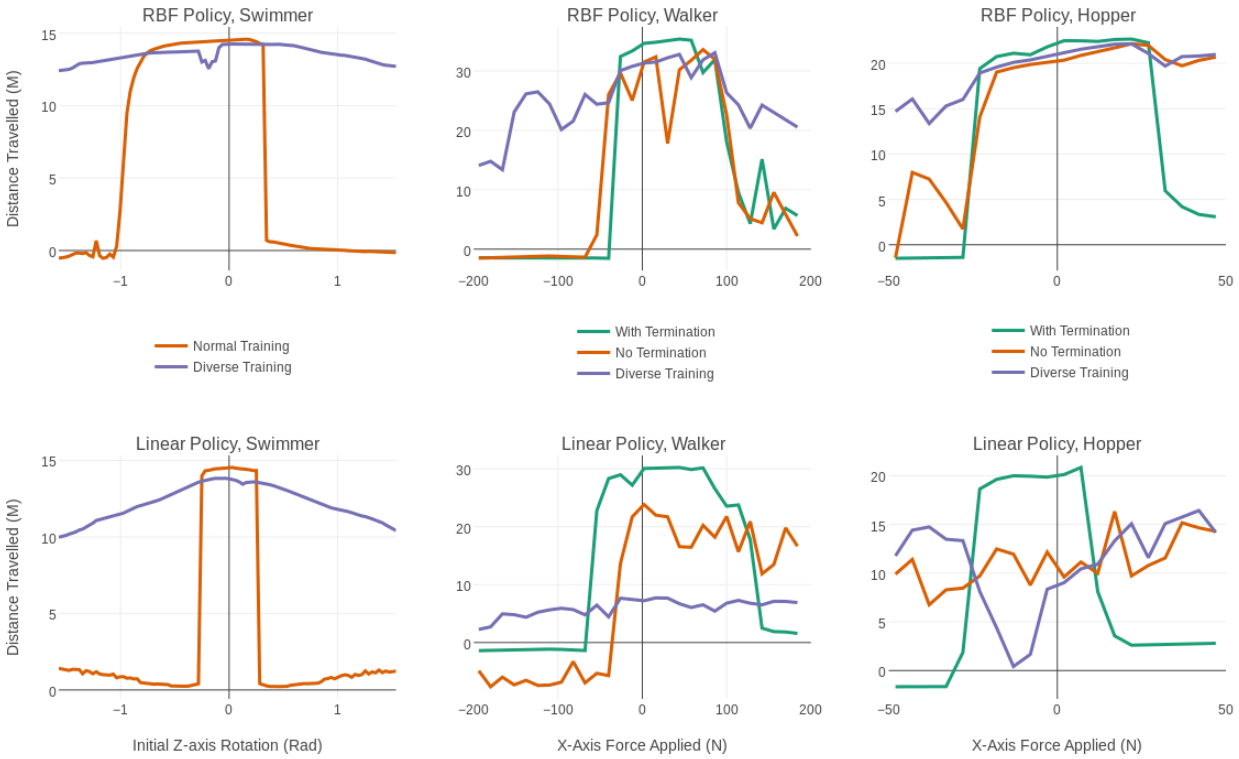

Figure 4. We test policy robustness by measuring distanced traveled in the swimmer, walker, and hopper tasks for three training configurations: with termination based on state, without termination but with peaked initial state distribution, and with diverse initialization. Swimmer does not have a termination option, so we consider only two configurations. For the case of swimmer, the perturbation is changing the heading angle between $-\pi/2.0$ and $\pi/2.0$, and in the case of walker and hopper, applying a force as the model is initialized for 0.5 seconds along its axis of movement. All models are initialized with the same positions and velocities. We can see that beyond some threshold, the 'with termination' training configuration fails to generalize-in most cases, the policy has given up and the system stops moving. Some interesting features of this test include the symmetry in the swimmer with linear policy, and the asymmetry in the RBF policy–the randomized weights may induce the swimmer to favor one side. Generally, RBF policies perform better than Linear policies. In the case of Walker, a Linear policy with diverse training is a very consistent performer, but is unable to achieve very high scores. We see that without diverse training, the Hopper with RBF policy never learns to recover from a backwards fall, but is quite capable with forward falls. It should be noted that in the Hopper with Linear policy and diverse training, the hopper remained in a stable upright position, choosing not to make forward progress when gently pertubed.

5.2. Towards Harder Tasks

One explanation for the good performance of policy gradient methods, even with simplistic policies, is that the considered locomotion tasks are conducive to beneficial random exploration. In both the original Gym-v1 tasks, as well as the modified tasks, the agent is not required to find a sequence of controls with very delayed reward like in the case of tasks like Mountain Car or Montezuma's Revenge.

We were unable to train a successful policy with the linear or RBF architecture on the Humanoid-v1 task. This possibly indicates that the Humanoid-v1 task poses representational challenges and likely requires directed exploration. Manipulation tasks, while not considered in this work, also pose similar challenges.

5.3. Robustness

The generalization and robustness we demonstrate in this work arise entirely out of the states the agent visits during training. No perturbations were applied during the training phase. Designing tasks with model ensembles (Rajeswaran et al., 2016) or perturbation forces during training could potentially further improve the robustness. We plan to pursue this in the near future.

5.4. Interactive Control

Interactive control tasks have been suggested and successfully solved in prior work (Tassa et al., 2012; Mordatch et al., 2015). Even though the interactive tasks that we study are not as complex as those in Mordatch et al. (2015),

to our knowledge, this is the first work that considers model-free approaches for continuous control tasks where a user can interact with the learned policy. As observed in prior work, interactive perturbations and ability to interactively set goals are an interesting way to study and characterize robustness and generalization as the human user can deliberately target policy weaknesses. Such approaches could also pave way for hierarchical planning, where the goals can be modulated by a higher level agent.

References

Al Borno, M., de Lasa, M., and Hertzmann, A. Trajectory Optimization for Full-Body Movements with Complex Contacts. *IEEE Transactions on Visualization and Computer Graphics*, 2013.

Amari, S. Natural gradient works efficiently in learning. *Neural Computation*, 10:251–276, 1998.

Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. OpenAI Gym, 2016.

Duan, Y., Chen, X., Houthooft, R., Schulman, J., and Abbeel, P. Benchmarking deep reinforcement learning for continuous control. In *ICML*, 2016.

Erez, Tom, Tassa, Yuval, and Todorov, Emanuel. Infinite-horizon model predictive control for periodic tasks with contacts. In *RSS*, 2011.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.

Kakade, S. A natural policy gradient. In *NIPS*, 2001.

Lillicrap, T., Hunt, J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. *ArXiv e-prints*, September 2015.

Mnih, V. et al. Human-level control through deep reinforcement learning. *Nature*, 518, 2015.

Mordatch, I., Todorov, E., and Popovic, Z. Discovery of complex behaviors through contact-invariant optimization. *ACM SIGGRAPH*, 2012.

Mordatch, I., Lowrey, K., Andrew, G., Popovic, Z., and Todorov, E. Interactive Control of Diverse Complex Characters with Neural Networks. In *NIPS*, 2015.

Rahimi, A. and Recht, B. Random Features for Large-Scale Kernel Machines. In *NIPS*, 2007.

Rajeswaran, A., Ghotra, S., Ravindran, B., and Levine, S. EPOpt: Learning Robust Neural Network Policies Using Model Ensembles. *ArXiv e-prints*, 2016.

Schulman, J., Levine, S., Moritz, P., Jordan, M., and Abbeel, P. Trust region policy optimization. In *ICML*, 2015.

Schulman, J., Moritz, P., Levine, S., Jordan, M., and Abbeel, P. High-dimensional continuous control using generalized advantage estimation. In *ICLR*, 2016.

Silver, D. et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529, 2016.

Tassa, Y., Erez, T., and Todorov, E. Synthesis and stabilization of complex behaviors through online trajectory optimization. *International Conference on Intelligent Robots and Systems*, 2012.

Todorov, E., Erez, T., and Tassa, Y. MuJoCo: A physics engine for model-based control. In *International Conference on Intelligent Robots and Systems*, 2012.